

Acyclic Games and Iterative Voting *

Reshef Meir¹

Maria Polukarov²

Jeffrey S. Rosenschein³

Nicholas R. Jennings⁴

¹Technion—Israel Institute of Technology reshefm@ie.technion.ac.il

²University of Southampton, United Kingdom mp3@ecs.soton.ac.uk

³The Hebrew University of Jerusalem, Israel jeff@cs.huji.ac.il

⁴Imperial College, United Kingdom n.jennings@imperial.ac.uk

Abstract

We consider iterative voting models and position them within the general framework of acyclic games and game forms. More specifically, we classify convergence results based on the underlying assumptions on the agent scheduler (the order of players) and the action scheduler (which better-reply is played).

Our main technical result is providing a complete picture of conditions for acyclicity in several variations of Plurality voting. In particular, we show that (a) under the traditional lexicographic tie-breaking, the game converges for any order of players under a weak restriction on voters' actions; and (b) Plurality with randomized tie-breaking is not guaranteed to converge under arbitrary agent schedulers, but from any initial state there is *some* path of better-replies to a Nash equilibrium. We thus show a first separation between restricted-acyclicity and weak-acyclicity of game forms, thereby settling an open question from [Kukushkin, 2011]. In addition, we refute another conjecture regarding strongly-acyclic voting rules.

1 Introduction

Researchers in economics and game theory since Cournot [1838] had been developing a formal framework to study questions about acyclicity and convergence of local improvement dynamics in games.

Intuitively put, strong-acyclicity means that the game will converge regardless of the order of players/voters and how they select their action (as long as the moving agents are improving their utility in every step), i.e. that there are no cycles of better-replies whatsoever; Weak-acyclicity means that while cycles may occur, from any initial state (voting profile) there is at least one path of better-replies that leads to a Nash equilibrium; Restricted-acyclicity is a middle ground, requiring convergence for any order of players (agent scheduler), but allowing the action scheduler to restrict the way they choose among several available replies (e.g., only allowing best-replies). Most relevant to us is the work of Kukushkin [1999; 2002; 2011], who studied general characterizations of game forms that guarantee various notions of acyclicity.

A more recent field is *iterative voting*. In the iterative voting model, voters have fixed preferences and start from some announcement (e.g., sincerely report their preferences). Votes are aggregated via some predefined rule (e.g. Plurality), but can change their votes after observing the current announcements and outcome. The game proceeds in turns, where a single voter changes his vote at each turn, until no voter has objections and the final outcome is announced. This process is similar to online polls via Doodle or Facebook, where users can log-in at any time and change their vote. Similarly, in offline committees the participants can sometimes ask to change their vote, seeing the current outcome.

The formal study of iterative voting rules was initiated about 6 years ago in a AAAI paper that was a preliminary version of this one [Meir *et al.*, 2010]. Iterative voting papers typically focus on common voting rules

*A preliminary version of this paper has been presented at AAAI-2010 [Meir *et al.*, 2010].

such as Plurality and Borda, and study the conditions under which convergence of the iterative process to a Nash equilibrium is guaranteed. Most results in the field consider best-reply dynamics [Lev and Rosenschein, 2012; Reyhani and Wilson, 2012; Obratzsova *et al.*, 2015].

While voting rules and game forms are essentially the same thing, the iterative voting literature has remained largely detached from the more general literature on acyclicity in games. Bridging this gap is the main conceptual contribution of this work, for two reasons. First, understanding the conditions that entail acyclicity of games and game forms is crucial to the understanding of iterative voting scenarios, and to properly compare convergence results (e.g. convergence of best-reply dynamics is a special case of restricted acyclicity). Likewise, convergence results for specific voting rules under best/better-reply dynamics may shed light on more general questions regarding acyclicity. Building on the formalism of Kukushkin [2011] for strong/ restricted/ weak-acyclicity of game forms, we re-interpret in this paper both known and new results on convergence of better- and best-reply in voting games, and answer some open questions.

1.1 Related work

Kukushkin [2011] provided several partial characterizations for game forms with strong acyclicity. In particular, he showed that if we further strengthen the acyclicity requirement to demand an ordinal potential, then this is attained if and only if the game form is dictatorial, i.e., there is at most one voter that can affect the outcome. He further characterized game forms that are strongly acyclic under *coalitional improvements*, and provided broad classes of game forms that are “almost unrestricted acyclic,” i.e. restricted-acyclic under mild restrictions on voters’ actions. Other partial characterizations have been provided for acyclicity in complete information extensive-form games [Boros *et al.*, 2008; Andersson *et al.*, 2010]. Some of this work is explained in more detail in the following sections.

The study of classes of games (i.e. game forms with utilities) that are guaranteed to be acyclic or weakly acyclic attracted much attention, in particular regarding the existence and properties of potential functions [Monderer and Shapley, 1996; Milchtaich, 1996; Fabrikant *et al.*, 2010; Apt and Simon, 2012].

Strategic voting The notion of strategic voting has been highlighted in research on Social Choice as crucial to understanding the relationship between preferences of a population and the final outcome of elections. In various applications (ranging from political domains to artificial intelligence [AI]), the most widely used voting rule is Plurality, in which each voter has one vote and the winner is the candidate who received the highest number of votes. While it is known that no reasonable voting rule is completely immune to strategic behavior [Gibbard, 1973; Satterthwaite, 1975], Plurality has been shown to be particularly susceptible, both in theory [Saari, 1990; Friedgut *et al.*, 2011] and in practice [Forsythe *et al.*, 1996]. This makes the analysis of any election campaign—even one where the simple Plurality rule is used—a challenging task. As voters may speculate and counter-speculate, it would be beneficial to have formal tools that would help us understand (and perhaps predict) the final outcome.

In particular, natural tools for this task include the well-studied solution concepts developed for normal form games, such as better/best responses, dominant strategies or different variants of equilibrium. Now, while voting settings are not commonly presented in this way, several natural formulations have been proposed in the past [Dhillon and Lockwood, 2004; Chopra *et al.*, 2004; Sertel and Sanver, 2004; Falik *et al.*, 2012; Messner and Polborn, 2002]. These formulations are extremely simple for Plurality voting games, where voters have only a few available ways to vote. Specifically, some of this previous work has been devoted to the analysis of solution concepts such as *elimination of dominated strategies* [Dhillon and Lockwood, 2004] and *strong equilibria* [Sertel and Sanver, 2004]. There has been other multi-step voting procedures that have been proposed in the literature, such as iterated majority vote [Airiau and Endriss, 2009] and extensive form games where voters vote one by one [Desmedt and Elkind, 2010]. In contrast to iterative voting, these models are inconsistent with the better-reply dynamics in normal form games, and are analyzed via different techniques. A model more similar to ours was recently studied in [Elkind *et al.*, 2015], where voters can choose between voting truthfully, and manipulating under the assumption that everyone else are truthful.

Convergence of better-reply dynamics in iterative voting for particular voting rules has been studied extensively in the computational social choice literature. We summarize and compare these findings with ours in the concluding section, and in particular in Table 1.

An important question in every model of strategic voting, including iterative voting, is whether the reached equilibrium is good for the society according to various metrics. Branzei et al.[2013] showed bounds on the *dynamic price of anarchy*, i.e. how far can the final outcome be from the initial truthful outcome. Other work used simulations to show that iterative voting may improve the social welfare or Condorcet efficiency [Grandi et al., 2013; Meir et al., 2014; Koolyk et al., 2016], but typically under the assumptions that voters use various heuristics.

Biased and sophisticated voting Some recent work on iterative voting deals with voters who are uncertain, truth-biased, lazy-biased, bounded-rational, non-myopic, or apply some other restrictions and/or heuristics that diverge from the standard notion of better-reply in games [Reijngoud and Endriss, 2012; Gohar, 2012; Grandi et al., 2013; Obratzsova et al., 2013; Meir et al., 2014; Rabinovich et al., 2015; Obratzsova et al., 2015; Meir, 2015]. Although the framework is suitable for studying such iterative dynamics as well, this paper deals exclusively with myopic better-reply dynamics.¹

1.2 Contribution and structure

The paper unfolds as follows. In Section 2, we define the iterative voting model within the more general framework of game forms and acyclicity properties. In Section 3 we consider strong acyclicity, and settle an open question regarding the existence of acyclic non-separable game forms. Section 4 focuses on order-free acyclicity of the Plurality rule. Our main result in this section shows that to guarantee convergence, it is necessary and sufficient that voters restrict their actions in a natural way that we term *direct reply*—meaning that a voter will only reassign his vote to a candidate that will become a winner as a result. In Section 5, we use variations of Plurality to show a strict separation between restricted acyclicity and weak acyclicity, thereby settling another open question. We conclude in Section 6.

2 Preliminaries

We usually denote sets by uppercase letters (e.g., A, B, \dots), and vectors by bold letters (e.g., $\mathbf{a} = (a_1, \dots, a_n)$).

2.1 Voting rules and game forms

There is a set C of m alternatives (or *candidates*), and a set N of n strategic agents, or *voters*. A game form (also called a *voting rule*) f allows each agent $i \in N$ to select an action a_i from a set of messages A_i . Thus the input to f is a vector $\mathbf{a} = (a_1, \dots, a_n)$ called an *action profile*. We also refer to a_i as the *vote* of agent i in profile \mathbf{a} . Then, f chooses a winning alternative—i.e., it is a function $f : \mathcal{A} \rightarrow C$, where $\mathcal{A} = \times_{i \in N} A_i$. See Fig. 1 for examples.

A voting rule f is *standard* if $A_i = A$ for all i , and A is either $\pi(C)$ (the set of permutations over C) or a coarsening of $\pi(C)$. Thus most common voting rules except Approval are standard. Mixed strategies are not allowed. The definitions in this section apply to all voting rules unless stated otherwise. For a permutation $P \in \pi(C)$, We denote by $\text{top}(P)$ the first element in P .

Plurality In the Plurality voting rule we have that $A = C$, and the winner is the candidate with the most votes. We allow for a broader set of “Plurality game forms” by considering both weighted and fixed voters, and varying the tie-breaking method. Each of the strategic voters $i \in N$ has an integer weight $w_i \in \mathbb{N}$. In addition, there are \hat{n} “fixed voters” who do not play strategically or change their vote. The vector $\hat{\mathbf{s}} \in \mathbb{N}^m$ (called “initial score vector”) specifies the number of fixed votes for each candidate. Weights and initial scores are part of the game form.²

¹We do consider however two standard ways to handle ties that slightly relax the better-reply definition. See Section 4.4.

²All of our results still hold if there are no fixed voters, but allowing fixed voters enables the introduction of simpler examples, and facilitates some of the proofs, see Remark 4.1. For further discussion on fixed voters see [Elkind et al., 2015].

f_1	a	b	c	f_2	a	b	c	f_3	x	y	f_4	x	y	z	w
a	a	a	a	a	a	a	a	a	a	b	a	ax	ay	az	aw
b	b	b	b	b	a	b	b	b	b	c	b	bx	by	bz	bw
c	c	c	c	c	a	b	c	c	c	a	c	cx	cy	cz	cw

Figure 1: Four examples of game forms with two agents. f_1 is a dictatorial game form with 3 candidates (the row agent is the dictator). f_2 is the Plurality voting rule with 3 candidates and lexicographic tie-breaking. f_3 and f_4 are non-standard game forms. In f_3 , $A_1 = C = \{a, b, c\}$, $A_2 = \{x, y\}$. Note that f_4 is completely general (there are 3×4 possible outcomes in C , one for each voting profile) and can represent any 3-by-4 game.

$f_{\mathbf{w}, \hat{\mathbf{s}}}^{PL}$	a	b	c
a	(14, 9, 3) { a }	(10, 13, 3) { b }	(10, 9, 7) { a }
b	(11, 12, 3) { b }	(7, 16, 3) { b }	(7, 12, 7) { b }
c	(11, 9, 6) { a }	(7, 13, 6) { b }	(7, 9, 10) { c }

Figure 2: A game form $f_{\mathbf{w}, \hat{\mathbf{s}}}^{PL}$, where $N = \{1, 2\}$, $A_1 = A_2 = C = \{a, b, c\}$, $\hat{\mathbf{s}} = (7, 9, 3)$ and $\mathbf{w} = (3, 4)$ (i.e., voter 1 has weight 3 and voter 2 has weight 4). The table shows the final score vector $\mathbf{s}_{(a_1, a_2)}$ for every joint action of the two voters, and the respective winning candidate $f_{\mathbf{w}, \hat{\mathbf{s}}}^{PL}(a_1, a_2)$ in curly brackets.

The *final score* of c for a given profile $\mathbf{a} \in A^n$ in the Plurality game form $f_{\mathbf{w}, \hat{\mathbf{s}}}$ is the total weight of voters that vote c . We denote the final score vector by $\mathbf{s}_{\hat{\mathbf{s}}, \mathbf{w}, \mathbf{a}}$ (often just $\mathbf{s}_{\mathbf{a}}$ or \mathbf{s} when the other parameters are clear from the context), where $s(c) = \hat{s}(c) + \sum_{i \in N: a_i = c} w_i$.

Thus the Plurality rule selects some candidate from $W = \operatorname{argmax}_{c \in C} s_{\hat{\mathbf{s}}, \mathbf{w}, \mathbf{a}}(c)$, breaking ties according to some specified method. The two primary variations we consider are $f_{\hat{\mathbf{s}}, \mathbf{w}}^{PL}$ which breaks ties lexicographically, and $f_{\hat{\mathbf{s}}, \mathbf{w}}^{PR}$ which selects a winner from W uniformly at random. As with \mathbf{s} , we omit the scripts \mathbf{w} and $\hat{\mathbf{s}}$ when they are clear from the context.

For illustration, consider an example in Fig. 2, demonstrating a specific weighted Plurality game form with two agents.

2.2 Incentives

Games are attained by adding either cardinal or ordinal utility to a game form. The linear order relation $Q_i \in \pi(C)$ reflects the preferences of agent i . That is, i prefers c over c' (denoted $c \succ_i c'$) if $(c, c') \in Q_i$. The vector containing the preferences of all n agents is called a *preference profile*, and is denoted by $\mathbf{Q} = (Q_1, \dots, Q_n)$. The game form f , coupled with a preference profile \mathbf{Q} , defines an ordinal utility normal form game $G = \langle f, \mathbf{Q} \rangle$ with n agents, where agent i prefers outcome $f(\mathbf{a})$ over outcome $f(\mathbf{a}')$ if $f(\mathbf{a}) \succ_i f(\mathbf{a}')$. In standard game forms the action a_i may indicate the agent's preferences, hence their common identification with voting rules.

Improvement steps and equilibria Having defined a normal form game, we can now apply standard solution concepts. Let $G = \langle f, \mathbf{Q} \rangle$ be a game, and let $\mathbf{a} = (\mathbf{a}_{-i}, a_i)$ be a joint action in G .

We denote by $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ an *individual improvement step*, if (1) \mathbf{a}, \mathbf{a}' differ only by the action of player i ; and (2) $f(\mathbf{a}_{-i}, a'_i) \succ_i f(\mathbf{a}_{-i}, a_i)$. We sometimes omit the actions of the other voters \mathbf{a}_{-i} when they are clear from the context, only writing $a_i \xrightarrow{i} a'_i$. We denote by $I_i(\mathbf{a}) \subseteq A_i$ the set of actions a'_i s.t. $a_i \xrightarrow{i} a'_i$ is an improvement step of agent i in \mathbf{a} , and $I(\mathbf{a}) = \bigcup_{i \in N} \bigcup_{a'_i \in I_i(\mathbf{a})} (\mathbf{a}_{-i}, a'_i)$. $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is called a *best reply* if a'_i is i 's most preferred candidate in $I_i(\mathbf{a})$.

A joint action \mathbf{a} is a (pure) *Nash equilibrium* (NE) in G if $I(\mathbf{a}) = \emptyset$. That is, no agent can gain by changing his vote, provided that others keep their strategies unchanged. A priori, a game with pure strategies does not have to admit any NE.

$\langle f, \mathbf{Q}^1 \rangle$	a	b	* c
* a	{ a } 3, 2	{b} 2, 1	* { a } 3, 2
b	{b} 2, 1	{ b } 2, 1	{b} 2, 1
c	{a} 3, 2	{b} 2, 1	{c} 1, 3

Figure 3: A game $G = \langle f, \mathbf{Q}^1 \rangle$, where $f = f_{\mathbf{w}, \mathbf{s}}^{PL}$ is as in Fig. 2, and \mathbf{Q}^1 is defined by $a \succ_1 b \succ_1 c$ and $c \succ_2 a \succ_2 b$. The table shows the ordinal utility of the outcome to each agent, where 3 means the best candidate. **Bold** outcomes are the NE points. Here the truthful vote (marked with *) is also a NE.

$\langle f, \mathbf{Q}^2 \rangle$	a	b	* c
* a	{a} 3, 1	{ b } 1, 2	* {a} 3, 1
b	{b} 1, 2	{ b } 1, 2	{b} 1, 2
c	{a} 3, 1	{b} 1, 2	{c} 2, 3

Figure 4: This game has the same game form as in Fig. 2, and the preference profile \mathbf{Q}^2 is $a \succ_1 c \succ_1 b$ and $c \succ_2 b \succ_2 a$. In this case, the truthful vote $\mathbf{a}^*(\mathbf{Q}^2)$ is not a NE.

Now, observe that when f is a standard voting rule the preference profile \mathbf{Q} induces a special joint action $\mathbf{a}^* = \mathbf{a}^*(\mathbf{Q})$, termed the *truthful state*, where a_i^* equals (the coarsening of) Q_i . E.g. in Plurality $a_i^* = \text{top}(Q_i)$. We refer to $f(\mathbf{a}^*)$ as the *truthful outcome* of the game $\langle f, \mathbf{Q} \rangle$.

The truthful state may or may not be included in the NE points of the game, as can be seen from Tables 3 and 4 that demonstrate games that are induced by adding incentives to the game form shown in Fig. 2, and indicate the truthful states and the NE points in these games.

2.3 Iterative Games

We consider natural *dynamics* in iterative games. Assume that agents start by announcing some initial profile \mathbf{a}^0 , and then proceed as follows: at each step t a single agent i may change his vote to $a_i^t \in I_i(\mathbf{a}^{t-1})$, resulting in a new state (joint action) $\mathbf{a}^t = (\mathbf{a}_{-i}^{t-1}, a_i^t)$. The process ends when no agent has objections, and the outcome is set by the last state.

Local improvement graphs and schedulers Any game G induces a directed graph whose vertices are all action profiles (states) \mathcal{A} , and edges are all local improvement steps [Young, 1993; Andersson *et al.*, 2010]. The pure Nash equilibria of G are all states with no outgoing edges. Since a state may have multiple outgoing edges ($|I(\mathbf{a})| > 1$), we need to specify which one is selected in a given play.

A *scheduler* ϕ selects which edge is followed at state \mathbf{a} at any step of the game [Apt and Simon, 2012]. The scheduler can be decomposed into two parts, namely selecting an agent i to play (agent scheduler ϕ^N), and selecting an action in $I_i(\mathbf{a})$ (action scheduler ϕ^A), where $\phi = (\phi^N, \phi^A)$. We note that a scheduler may or may not depend on the history or other factors, but this does not affect any of our results.

Convergence and acyclicity Given a game G , an initial action profile \mathbf{a}^0 and a scheduler ϕ , we get a unique (possibly infinite) path of steps.³ Also, it is immediate to see that the path is finite if and only if it reaches a Nash equilibrium (which is the last state in the path). We say that the triple $\langle G, \mathbf{a}^0, \phi \rangle$ *converges* if the induced path is finite.

Following [Monderer and Shapley, 1996; Milchtaich, 1996], a game G has the *finite individual improvement property* (we say that G is FIP), if $\langle G, \mathbf{a}^0, \phi \rangle$ converges for *any* \mathbf{a}^0 and scheduler ϕ . Games that are FIP are also known as *acyclic games* and as *generalized ordinal potential games* [Monderer and Shapley, 1996].

It is quite easy to see that not all Plurality games are FIP (see examples in Section 4). However, there are alternative, weaker notions of acyclicity and convergence.

³By “step” we mean an individual improvement step, unless specified otherwise.

- A game G is *weakly-FIP* if there is *some* scheduler ϕ such that $\langle G, \mathbf{a}^0, \phi \rangle$ converges for any \mathbf{a}^0 . Such games are known as *weakly acyclic*, or as ϕ -potential games [Apt and Simon, 2012].
- A game G is *restricted-FIP* if there is *some action scheduler* ϕ^A such that $\langle G, \mathbf{a}^0, (\phi^N, \phi^A) \rangle$ converges for any \mathbf{a}^0 and ϕ^N [Kukushkin, 2011]. We term such games as *order-free acyclic*.


Intuitively, restricted FIP means that there is some restriction players can adopt s.t. convergence is guaranteed regardless of the order in which they play. Kukushkin identifies a particular restriction of interest, namely restriction to best-reply improvements, and defines the *finite best-reply property* (FBRP) and its weak and restricted analogs. We emphasize that an action scheduler *must* select an action in $I_i(\mathbf{a})$, if one exists. Thus restricted dynamics that may disallow all available actions (as in [Gohar, 2012; Grandi *et al.*, 2013]) do not fall under the definition of restricted-FIP (but can be considered as separate dynamics).

We identify a different restriction, namely *direct reply*, that is well defined under the Plurality rule. Formally, a step $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is a direct reply if $f(\mathbf{a}') = a'_i$, i.e., if i votes for the new winner (see labeled examples in Section 4). Another rule where a natural direct reply exists is Veto, where a voter can veto the current winner [Lev and Rosenschein, 2012].

ϕ^A is direct if it always selects a direct reply. We get the following definitions for a Plurality game G , where FDRP stands for *finite direct reply property*:

- G is *FDRP* if $\langle G, \mathbf{a}^0, \phi \rangle$ converges for any \mathbf{a}^0 and any direct ϕ .
- G is *weakly-FDRP* if there is a direct ϕ such that $\langle G, \mathbf{a}^0, \phi \rangle$ converges for any \mathbf{a}^0 .
- G is *restricted-FDRP* if there is a direct ϕ^A such that $\langle G, \mathbf{a}^0, (\phi^N, \phi^A) \rangle$ converges for any \mathbf{a}^0 and ϕ^N .
- FDBRP means that replies are both best and direct. Note that it is unique and thus cannot be further restricted.

Finally, a game form f has the X property (where X is any of the above versions of finite improvement) if $\langle f, \mathbf{Q} \rangle$ is X for all preference profiles $\mathbf{Q} \in (\pi(C))^n$. We have the following entailments, both for games and for game forms. The third row is only relevant for Plurality/Veto.



Kukushkin notes that there are no known examples of game forms that are weak-FIP, but not restricted-FIP. We settle this question later in Section 5.2.

Convergence from the truth We say that a game G is *FIP from state* \mathbf{a} if $\langle G, \mathbf{a}, \phi \rangle$ converges for any ϕ . Clearly a game is FIP iff it is FIP from \mathbf{a} for any $\mathbf{a} \in A^n$. The definitions for other all other notions of finite improvement properties are analogous.

We are particularly interested in convergence from the truthful state \mathbf{a}^* . This is since: a. it is rather plausible to assume that agents will start by voting truthfully, especially when not sure about others' preferences; and b. even with complete information, they may be inclined to start truthfully, as they can always later change their vote.

Heuristic voting Much work on iterative voting deals with heuristics, rather than best- or better-replies. Strong, Restricted, and Weak convergence properties can be defined the same way, where the only difference is the way we define $I_i(\mathbf{a})$ (i.e., all steps that are allowed for agent i at state \mathbf{a} by the considered heuristics). For example, *truth-bias* assumes that if a voter does not have any local improvement step, she reverts to her truthful vote [Obraztsova *et al.*, 2013]. Some heuristics are already restricted to a single action (for example, “k-pragmatist” [Grandi *et al.*, 2013]). In these cases the only meaningful distinction is between FIP and weak-FIP. In this paper we do not consider heuristic voting.

3 Strong Acyclicity

An *ordinal potential* is a function that strictly increases if and only if some agent plays a better-reply [Monderer and Shapley, 1996]. A *generalized ordinal potential* is a function that strictly increases with every better-reply, but may also increase with other steps. Clearly, a game is FIP if and only if it has a generalized potential (by a topological sort of the better-reply graph).

Theorem 1 (Kukushkin [2011]). *A game form f guarantees an ordinal potential (i.e. every derived game has an ordinal potential function) if and only if f is a dictatorship.*

We emphasize that this result *does not* preclude the existence of other game forms with FIP (*generalized* ordinal potential). Indeed, Kukushkin provides a partial characterization of FIP game forms. For example, a rule where there is a linear order L over C , and the winner is the first candidate according to L that is top-ranked by at least one voter.

A game form f is called “separable” [Kukushkin, 2011] if there are mappings $g_i : A_i \rightarrow C$ for $i \in N$ s.t. for all $\mathbf{a} \in \mathcal{A}$, $f(\mathbf{a}) \in \{g_1(a_1), g_2(a_2), \dots, g_n(a_n)\}$. That is, the vote of each voter is mapped to a single candidate via some function g_i , and the outcome is always one of the candidates in the range. Examples of separable rules include Plurality and dictatorial rules, in both of which g_i are the identity functions.

Conjecture 2 (Kukushkin [2011]). *Any FIP game form is separable.*

Some weaker variations of this conjecture have been proved. In particular, for game forms with finite *coalitional improvement* property [Kukushkin, 2011], and for FIP game forms with $n = 2$ voters [Boros *et al.*, 2010] (separable game forms are called “assignable” there). We next show that for sufficiently large n , there are non-separable FIP game forms, thereby refuting the conjecture. Our proof uses the probabilistic method: we sample a game form from some space, and prove that with positive probability it must be non-separable and FIP.

Theorem 3. *For any $n \geq 20$, there is a non-separable game form f_n s.t. f_n is FIP.*

Proof. Let $C = \{a^1, \dots, a^{2n}\} \cup \{z\}$. Let $A_i = \{x, y\}$ for each voter. Thus f_n is a function from the n dimensional binary cube $\mathcal{B} = \{x, y\}^n$ to C . We select $2n$ profiles $\mathbf{a}^1, \dots, \mathbf{a}^{2n}$ uniformly at random, i.i.d. from \mathcal{B} (allowing repetitions), and define $f_n(\mathbf{a}^j) = a^j$ for all $j \in \{1, \dots, 2n\}$. For all other $2^n - 2n$ profiles we define $f_n(\mathbf{a}) = z$.

For any two profiles \mathbf{a}, \mathbf{a}' , let $d(\mathbf{a}, \mathbf{a}')$ be the number of voters that disagree in \mathbf{a}, \mathbf{a}' (the Manhattan distance on the cube). Let $B \subseteq \mathcal{B}$ be all $2n$ profiles whose outcome is not z . For $j, j' \leq 2n$, denote by $p_{j,j'}$ the probability that $d(\mathbf{a}^j, \mathbf{a}^{j'}) \leq 2$, and by $X_{j,j'}$ the corresponding indicator random variable. Since both of $\mathbf{a}^j, \mathbf{a}^{j'}$ were sampled uniformly i.i.d., and there are less than n^2 profiles within distance 2 from \mathbf{a}^j , we get that $p_{j,j'} \leq \frac{n^2}{2^n}$.

Next, by the union bound,

$$Pr(\exists X_{j,j'} = 1) \leq \sum_{j \leq 2n, j' \leq 2n} Pr(X_{j,j'}) = \sum_{j \leq 2n, j' \leq 2n} p_{j,j'} \leq (2n)^2 \frac{n^2}{2^n} = \frac{4n^2}{2^n},$$

which is strictly less than 1 for $n \geq 20$. Thus w.p. > 0 we get $X_{j,j'} = 0$ for all j, j' . In particular there is at least one such game form f_n^* where $X_{j,j'} = 0$ for all j, j' . We argue that f_n^* is both FIP and non-separable.⁴

Assume towards a contradiction that there is some cycle of better-replies in f_n^* . Then there must be a path containing at least 3 distinct outcomes, and thus at least 2 profiles from B . Denote these profiles by \mathbf{a}, \mathbf{b} . Since $X_{j,j'} = 0$ for all j, j' , we have that any path between \mathbf{a} and \mathbf{b} is of length at least 3, and that the path must contain at least two consequent states whose outcome is z . This path cannot be a better-reply path, since a better reply must change the outcome. Hence we get a contradiction and f_n^* is FIP.

Finally, note that since $X_{j,j'} = 0$ for all j, j' , in particular \mathbf{a}^j are all distinct profiles, and thus f_n^* has $2n + 1 > \sum_{i \leq n} |A_i|$ possible outcomes. In contrast, for any separable rule f the size of the range of f is at most $\sum_{i \leq n} |A_i|$, since $f(\mathbf{a}) = g_i(b)$ for some $i \in N$ and $b \in A_i$. This means that f_n^* is non-separable. \square

⁴Using the Hamming error-correcting code [Hamming, 1950], it is in fact possible to explicitly construct f_n^* for as few as $n = 7$ voters. The rest of the proof remains the same.

For most common voting rules, separable or not, it is easy to find examples where some cycles occur. Thus one should focus on the weaker notions of convergence discussed in Section 1, which is what we do in the remainder of the paper.

4 Order-Free Acyclicity: Plurality

Improvement steps in Plurality Recall that along a given path, $\mathbf{a}^t \in A^n = C^n$ denotes the voting profile at time t . We denote by $\mathbf{s}^t = \mathbf{s}_{\mathbf{a}^t}$ the score vector at time t ; by $cw^t = f^{PL}(\mathbf{a}^t)$ the candidate that wins at time t ; and by $sw^t = \check{\mathbf{s}}^t(cw^t)$ the highest score at time t (including tie-breaking if it applies).

Suppose that agent i has an improvement step (a.k.a. *better reply*) $a_i^{t-1} \xrightarrow{i} a_i^t$ at time t . We classify all possible steps into the following types (an example of such a step appears in parentheses):

Type 1. from $a_i^{t-1} \neq cw^{t-1}$ to $a_i^t = cw^t$; (step 1 in Ex.6a.)

Type 2. from $a_i^{t-1} = cw^{t-1}$ to $a_i^t = cw^t$; (step 1 in Ex.6b.),

Type 3. from $a_i^{t-1} = cw^{t-1}$ to $a_i^t \neq cw^t$; (step 2 in Ex.6a.)

Note that steps of type 1 and 2 are direct, whereas type 3 steps are indirect.

4.1 Lexicographic Tie-Breaking

In this section we assume that ties are broken lexicographically. Given some score vector \mathbf{s} , we denote by $\check{s}(c) \in \mathbb{R}$ the score of $c \in C$ that includes the lexicographic tie-breaking component. One way to formally define it is by setting $\check{s}(c) = s(c) + \frac{1}{m+1}(m - L(c))$, where $L(c)$ is the lexicographic index of candidate c . However the only important property of \check{s} is that $\check{s}(c) > \check{s}(c')$ if either $s(c) > s(c')$ or the score is equal and c has a higher priority (lower index) than c' .

Thus for Plurality with lexicographic tie-breaking, a given weight vector \mathbf{w} and a given initial score vector $\hat{\mathbf{s}}$, we denote the outcome by

$$f_{\hat{\mathbf{s}}, \mathbf{w}}^{PL}(\mathbf{a}) = \operatorname{argmax}_{c \in C} \check{s}_{\hat{\mathbf{s}}, \mathbf{w}, \mathbf{a}}(c).$$

As with \mathbf{s} , we omit the scripts \mathbf{w} , $\hat{\mathbf{s}}$ and PL when they are clear from the context.

Lemma 4. Consider a game $\langle f_{\mathbf{w}, \hat{\mathbf{s}}}^{PL}, \mathbf{Q} \rangle$. If there exists a better reply for a given agent i at state \mathbf{a}^{t-1} , then i has a direct best reply at state \mathbf{a}^{t-1} .

The proof is trivial under lexicographic tie-breaking, by letting i vote for her most preferred candidate among all better replies. In this case the direct best reply is also unique.

One implication of the lemma is that it is justified and natural to restrict our discussion to direct replies and focus on FDRP, as w.l.o.g. a voter always has a direct reply that is at least as good as any other reply.

Unweighted Voters Suppose all voters have unit weight. We start with our main result for this section.

Theorem 5. $f_{\hat{\mathbf{s}}}^{PL}$ is FDRP. Moreover, any path of direct replies will converge after at most $m^2 n^2$ steps. In particular, Plurality is order-free acyclic.

This extends a weaker version of the theorem that appeared in the preliminary version of this paper [Meir *et al.*, 2010], which only showed FDBRP. The bound on the number of direct-best-reply steps was recently improved to $O(mn)$ in [Reyhani and Wilson, 2012, Theorem 5.4].

Proof. By our restriction to direct replies, there can only be moves of types 1 and 2. We first consider moves of type 1, and inductively prove two invariants that yield a bound on the total number of such moves. Next, we bound the number of moves of type 2 by a given voter between any of his moves of type 1, which completes the proof.




Figure 5: An illustration of a type 1 move. Tie-breaking is in favor of the left most candidate.

Consider time $t - 1$ and denote the score of the current winner (including tie-breaking) by $\bar{s} = sw^{t-1}$. Suppose that a move $a \xrightarrow{i} b$ of type 1 occurs at time t : that is, $a \neq cw^{t-1}$ and $b = cw^t$. We then have (see Fig. 5):

$$\check{s}^t(b) = sw^t \geq sw^{t-1} = \bar{s} \geq \check{s}^{t-1}(a) = \check{s}^t(a) + 1. \quad (1)$$

We claim that at any later time $t' \geq t$ the following two invariants hold:

- I. Either there is a candidate $c \neq a$ whose score is at least $\bar{s} + 1$, or there are at least two candidates $c, c' \neq a$ whose score is at least \bar{s} . In particular it holds in either case that $sw^{t'} \geq \bar{s}$.
- II. The score of a does not increase: $\check{s}^{t'}(a) \leq \check{s}^t(a)$.

Note that this, coupled with Eq. (1), implies that candidate a will never win again, as its score will stay strictly below \bar{s} , and there will always be a candidate with a score of at least \bar{s} .

We now prove both invariants by induction on time t' . In the base case $t' = t$, (I) holds since both cw^{t-1} and b have a score of at least \bar{s} , and (II) holds trivially.

Assume by induction that both invariants hold until time $t' - 1$, and consider step t' by voter j . Due to (I), we either have at least two candidates whose score is at least \bar{s} , or a candidate with a score of at least $\bar{s} + 1$. Due to (II) and Eq. (1) we have that $\check{s}_{t'}(a) \leq \check{s}_t(a) < \bar{s} - 1$.

Let $d \xrightarrow{j} d'$ be the step at time t' by voter j (that is, $d = a_j^{t'-1}$, $d' = a_j^{t'}$). We first argue that $d' \neq a$: by adding the vote of j to a its score will still be strictly less than \bar{s} , whereas by removing a vote from any other candidate d , we still have at least one candidate c with score at least \bar{s} . Thus a cannot be a direct reply for any voter j , and (II) still holds after step t' .

It remains to show that (I) holds. If d is *not* one of the candidates in (I) with the score of at least \bar{s} at time $t' - 1$, then their score does not decrease after step t' , and we are done. Otherwise, we divide into the following cases:

1. At $t' - 1$, d is the (only) candidate with a score of at least $\bar{s} + 1$.
2. At $t' - 1$, candidates c, c' have scores of at least \bar{s} , and d is one of them (w.l.o.g. $d = c$).

In the first case, $\check{s}^{t'}(d) = \check{s}^{t-1}(d) - 1 \geq \bar{s} + 1 - 1 = \bar{s}$, whereas $\check{s}^{t'}(d') > \check{s}^{t'}(d) \geq \bar{s}$. Thus both d, d' have scores of at least \bar{s} at time t' , as required. In the second case, since only $c = d$ can lose votes, then if $d' \neq c'$,

$$\check{s}^{t'}(d') = sw^{t'} \geq \check{s}^{t'}(c') = \check{s}^{t'-1}(c') \geq \bar{s},$$

and thus both c', d' have scores of at least \bar{s} at time t , as required. If $d' = c'$, then and thus both c', d' have scores of at least \bar{s} at time t , as required. If $d' = c'$, then

$$\check{s}^{t'}(d') = \check{s}^{t'-1}(d') + 1 = \check{s}^{t'-1}(c') + 1 \geq \bar{s} + 1,$$

that is, d' has a score of at least $\bar{s} + 1$, as required.

Next, we demonstrate that invariants (I) and (II) supply us with a polynomial bound on the rate of convergence. Indeed, as we mentioned before, at every step of type 1, at least one candidate is ruled out permanently, and there are at most n times that a vote can be withdrawn from a given candidate. Also note that, since a type 2 move by a given voter i implies that he prefers a_i^t to a_i^{t-1} , each voter can make at most $m - 1$ type 2 moves before making a move of type 1. Hence, there are in total at most $m^2 n^2$ steps until convergence. \square

Furthermore, it is easy to show that if all voters start from the truthful state then type 2 moves never occur. Thus, the score of the winner never decreases, and the game converges in at most mn steps.

Next, we show that the restriction to direct replies is necessary to guarantee convergence, whereas a restriction to best replies is insufficient.

Proposition 6. f^{PL} is not FBRP, even from the truthful state. Moreover, there are: (a) a counterexample with two strategic agents and an arbitrary initial state; (b) a counterexample with three strategic agents and a truthful initial vote.

Remark 4.1. In this example and in others throughout the paper we use an initial score vector \hat{s} . However, this is w.l.o.g. since we could replace \hat{s} with additional voters that do not participate in the cycle. Initial scores are only useful to construct examples that are simpler and/or with fewer strategic agents. This holds for all negative results in the paper.⁵ For positive results, we have to show convergence for every initial scores \hat{s} .

Example 6a. $C = \{a, b, c\}$. We have a single fixed voter voting for a , thus $\hat{s} = (1, 0, 0)$. The preference profile is defined as $a \succ_1 b \succ_1 c$, $c \succ_2 b \succ_2 a$. The following cycle consists of better replies (the vector denotes the votes (a_1, a_2) at time t , the winner appears in curly brackets):

$$(b, c)\{a\} \xrightarrow{2} (b, b)\{b\} \xrightarrow{1} (c, b)\{a\} \xrightarrow{2} (c, c)\{c\} \xrightarrow{1} (b, c).$$

Note that all steps are best-replies, but the steps of agent 1 are indirect. \diamond

Example 6b. $C = \{a, b, c, d\}$. Candidates a, b , and c have 2 fixed voters each, thus $\hat{s} = (2, 2, 2, 0)$. We use 3 agents with the following preferences: $d \succ_1 a \succ_1 b \succ_1 c$, $c \succ_2 b \succ_2 a \succ_2 d$ and $d \succ_3 a \succ_3 b \succ_3 c$. Starting from the truthful state (d, c, d) the agents can make the following two improvement steps, which are direct best-replies (showing only the outcome s and the winner): $(2, 2, 3, 2)\{c\} \xrightarrow{1} (2, 3, 3, 1)\{b\} \xrightarrow{3} (3, 3, 3, 0)\{a\}$, after which agents 1 and 2 repeat the cycle shown in (6a). \diamond

Thus for the non-weighted lexicographic case Theorem 5 and Proposition 6 provide a clear-cut rule: direct replies guarantee convergence, whereas convergence is not guaranteed under other restrictions such as best reply or initial truthful vote. However, as the following section demonstrates, in the presence of weighted agents even direct replies may no longer converge.

Weighted Voters Next, we show that if the voters may have non-identical weights, then convergence to equilibrium is not guaranteed even if they start from the truthful state and use direct best replies.

Proposition 7. There is f_w^{PL} that is not restricted-FDRP, even from the truthful state.

Example 7. The initial fixed score of candidates $\{a, b, c, d\}$ is $\hat{s} = (0, 1, 2, 3)$. The weight of each voter $i \in \{1, 2, 3\}$ is i . The preference profile is as follows: $c \succ_1 d \succ_1 b \succ_1 a$, $b \succ_2 c \succ_2 a \succ_2 d$, and $a \succ_3 b \succ_3 c \succ_3 d$. We omit the rest of the proof. The initial truthful profile is thus $\mathbf{a}^0 = (c, b, a)$, which results in the score vector $\mathbf{s}^0 = (3, 3, 3, 3)$ where a is the winner.

$$\begin{array}{ccccc} \text{votes:} & (c, b, a) & \xrightarrow{1} & (d, b, a) & \xrightarrow{2} & (d, c, a) \\ \text{scores:} & (3, 3, 3, 3)\{a\} & & (3, 3, 2, 4)\{d\} & & (3, 1, 4, 4)\{c\} \\ & \uparrow_3 & & & & \downarrow_3 \\ & (c, b, b) & \xleftarrow{2} & (c, c, b) & \xleftarrow{1} & (d, c, b) \\ & (0, 6, 3, 3)\{b\} & & (0, 4, 5, 3)\{c\} & & (0, 4, 4, 4)\{b\} \end{array}$$

⁵Note that the remark does no longer hold if \hat{s} is used to construct a counter example for weak-FIP. However we use no such examples in this paper.

Our example shows a cycle of direct responses. Note that at every step there is only one direct reply available to the agent, thus it is not possible to eliminate the cycle by further restricting the action scheduler. \diamond

If there are *only two* weighted voters (and possibly other fixed voters), either restriction to direct reply or to a truthful initial state is sufficient to guarantee convergence.

Theorem 8. $f_{\hat{s},w}^{PL}$ is FDRP for $n = 2$.

Proof. Clearly, in one of the two first states, the agents vote for distinct candidates. At any later state, they must continue voting for distinct candidates, as every step must change the winner, and the other voter is always voting for the current winner. This means that the score of the winner strictly increases with every step (possibly except the first one). \square

Theorem 9. $f_{\hat{s},w}^{PL}$ is FIP from the truth for $n = 2$.

Proof. We show that the score of the winner can only increase. This clearly holds in the first step, which must be of type 1. Once again, we have that the two agents always vote for different candidates, and thus only steps that increase the score can change the identity of the winner. \square

Thus in either case convergence is guaranteed after at most $2m$ steps.

It remains an open question whether there is any restriction on better replies that guarantees order-free acyclicity in weighted games, i.e. if f_w^{PL} is restricted-FIP for $n > 2$. However Prop. 7 shows that if such restricted dynamic exists, it must make use of indirect replies, which is rather unnatural. We thus conjecture that such restricted dynamics does not exist.

4.2 Arbitrary tie-breaking

Lev and Rosenschein [2012] showed that for any positional scoring rule (including Plurality), we can assign some (deterministic) tie breaking rule, so that the resulting voting rule may contains cycles. For any positional scoring rule f_α with score vector α , denote by f_α^{LR} the same rule with the Lev-Rosenschein tie-breaking.

Proposition 10 (Theorem 1 in [Lev and Rosenschein, 2012]). f_α^{LR} is not FBRP for any α , even for $n = 2$, and even from the truth. In particular, Plurality with the Lev-Rosenschein tie-breaking (f^{PLR}) is not FBRP.

In fact, a slight modification of their example (switching a and b in voter 2's preferences) yields the following:

Proposition 11. f^{PLR} is not restricted-FIP, even for $n = 2$, and even from the truth.

4.3 Randomized tie-breaking

Compared to the previously considered deterministic rule, randomized tie-breaking has the advantage of being neutral—no specific candidate or voter is preferred over another. Formally, the game form $f_{\hat{s},w}^{PR}$ maps any state $\mathbf{a} \in A^n$ to the set $\text{argmax}_{c \in C} s_{\hat{s},w,\mathbf{a}}(c)$. Since under randomized tie-breaking there are multiple winners, let $W^t = f^{PR}(\mathbf{a}^t) \subseteq C$ denote the set of winners at time t .⁶ We define a direct reply $a_i^{t-1} \xrightarrow{i} a_i^t$ as one where $a_i^t \in W^t$.

If ties are broken randomly, \succ_i does *not* induce a complete order over outcomes. For instance, the order $a \succ_i b \succ_i c$ does not determine if i will prefer $\{b\}$ over $\{a, c\}$. However, we can naturally extend Q_i to a *partial preference order* over subsets. There are several standard extensions, using the following axioms:⁷

K (Kelly [Kelly, 1977]): (1) $(\forall a \in X, b \in W, a \succ_i b) \Rightarrow X \succ_i W$; (2) $(\forall a \in X, b \in W, a \succeq_i b) \Rightarrow X \succeq_i W$;

G (Gärdenfors [Gärdenfors, 1976]): $(\forall b \in W, a \succ_i b) \Rightarrow \{a\} \succ_i (\{a\} \cup W) \succ_i W$;

⁶This is a slight abuse of the notation we introduce in the beginning, where we defined the set of possible outcomes of f to be C . Here we allow any $W \in 2^C \setminus \{\emptyset\}$ as a possible outcome.

⁷We thank an anonymous reviewer for the references.

R (Responsiveness [Roth, 1985]): $a \succ_i b \iff \forall W \subseteq C \setminus \{a, b\}, (\{a\} \cup W) \succ_i (\{b\} \cup W)$.

The axioms reflect various beliefs a rational voter may have on the tie-breaking procedure: the K axiom reflects no assumptions whatsoever; The K+G axioms are consistent with tie-breaking according to a fixed and unknown order [Geist and Endriss, 2011]; and K+G+R axioms are consistent with random tie-breaking with equal probabilities (see Lemma 15 and Prop. 20). In this section we assume all axioms hold, however our results do not depend on these interpretations, and we do not specify the voter’s preferences in cases not covered by the above axioms. Under strict preferences, it also holds that G entails K [Endriss, 2013]. We can also define “weak” variants G2 and R2 for axioms G and R, by replacing all strict relations with weak ones, however as long as we restrict attention to strict preferences over elements the weak variants are not required.

For the following lemma we only need Axiom K, i.e. it does not depend on the voter’s tie-breaking assumptions.

Lemma 12. *If there exists a better-reply in $f_{\hat{s}, \mathbf{w}}^{PR}$ for agent i at state \mathbf{a}^{t-1} , then i has a direct best-reply.*

Proof. Suppose there is a better reply $a_i^{t-1} \xrightarrow{i} b$ at time $t - 1$. As some best reply always exists, denote by b' an arbitrary best reply. Let $W = f_{\hat{s}, \mathbf{w}}^{RP}(\mathbf{a}_{-i}^{t-1}, b')$, and let a' be the most preferred candidate of i in W . Then we argue that $a_i^{t-1} \xrightarrow{i} a'$ is a direct best reply of i . Since a' is a direct reply by definition, it is left to show that a' is a best reply (for the lexicographic case this follows immediately from $W = \{a'\}$ and $f^{PL}(\mathbf{a}_{-i}^{t-1}, a') = W = \{a'\}$).

If b' is a direct reply then $b' = a'$ and we are done. Thus assume that b' is not a direct reply from $a_i^{t-1} \xrightarrow{i}$. Then $b' \notin W$. By voting for $a' \in W$, we get that $f_{\hat{s}, \mathbf{w}}^{RP}(\mathbf{a}_{-i}^{t-1}, a') = \{a'\}$, i.e., a' remains the unique winner. If $|W| = 1$ then we are done as in the lexicographic case. Otherwise we apply Axiom K2 with $X = \{a'\}$, and get that $a' \succeq_i W$. That is,

$$f_{\hat{s}, \mathbf{w}}^{RP}(\mathbf{a}_{-i}^{t-1}, a') = \{a'\} \succeq_i W = f_{\hat{s}, \mathbf{w}}^{RP}(\mathbf{a}_{-i}^{t-1}, b'),$$

which means that a' is also a best-reply. \square

With weighted votes and random tie-breaking, there may not be any pure Nash equilibrium at all [Meir *et al.*, 2010]. We therefore restrict attention in the rest of this section to unweighted votes.

Proposition 13. *f^{PR} is not FIP.*

Example 13. $C = \{a, b, c\}$ with initial score $\hat{s} = (0, 1, 0)$. The initial state is $\mathbf{a}_0 = (a, a, b)$ —that is, $s(\mathbf{a}_0) = (2, 2, 0)$ and the outcome is the winner set $\{a, b\}$. The preferences are $a \succ_1 c \succ_1 b, b \succ_2 a \succ_2 c$ and $c \succ_3 b \succ_3 a$. We get the following cyclic sequence:

$$\begin{array}{ccccc} (2, 2, 0)\{a, b\} & \xrightarrow{2} & (1, 2, 1)\{b\} & \xrightarrow{1} & (0, 2, 2)\{b, c\} \\ \uparrow 3 & & & & \downarrow 3 \\ (1, 2, 1)\{b\} & \xleftarrow{1} & (2, 1, 1)\{a\} & \xleftarrow{2} & (1, 1, 2)\{c\} \end{array}$$

We emphasize that each step is justified as a better reply by either Axiom K or Axiom G. E.g, in the step of agent 2 in the top row, agent 2 prefers $b \succ_2 a$, and thus $b \succ_2 \{a, b\}$ by Axiom G. This will be used later in Section 4.4. \diamond

Theorem 14. *$f_{\hat{s}}^{PR}$ is FBRP from the truth.*

Proof. We denote the sets of winners and runnerups at time t as $W^t = f^{RP}(\mathbf{a}^t); R^t = \{c : s^t(c) = sw^t - 1\}$. We will show by induction that at any step $\mathbf{a}^{t-1} \xrightarrow{i} \mathbf{a}^t$:

1. $W^t \cup R^t \subseteq W^{t-1} \cup R^{t-1}$ (i.e., candidates not in $W^t \cup R^t$ will not be selected by any agent at a later time).
2. a_i^t is the most preferred candidate for i in $W^t \cup R^t$ (in particular, a best reply is a direct reply).
3. $a_i^{t-1} \succ_i a_i^t$ (in the terminology of [Meir *et al.*, 2014], this is a *compromise step*).

Since each voter can make at most $m - 1$ compromise steps, convergence is guaranteed within nm steps.

Assume that for some $t \geq 1$, all of the above holds for any $t' < t$ (so we prove the base case together with the other cases). Since \mathbf{a}^0 is truthful, the first step of any voter is always a compromise move. If i had already moved at some previous time $t' < t$, then $a_i^{t'}$ is most preferred in $W^{t'} \cup R^{t'}$.

By induction, $a = a_i^{t-1}$ is the most preferred candidate in some C' that contains $W^{t-1} \cup R^{t-1}$ ($C' = C$ in i 's first step, and $C' = W^{t'} \cup R^{t'}$ at any other step). Let x and y be i 's most preferred candidates in W^{t-1} and in R^{t-1} , respectively, and denote the best reply by $a' = a_i^t$. Each of a or a' may belong to W^{t-1} , to R^{t-1} , or to neither set. This means there are $3 \times 3 = 9$ cases to check. Fortunately, we can show that some of these cases immediately lead to a contradiction, and in the other cases all invariants 1-3 will hold after step t .

Consider first the case $a \in W^{t-1}$. Since a is most preferred in C' , it is strictly more preferred than any other candidate in W^{t-1} or in R^{t-1} (i.e., $a = x$). Thus if $a' \in W^{t-1}$ we get $W^t = \{a'\} \prec_i W^{t-1}$ by Axiom G. If $a' \in R^{t-1}$ we get $W^t = (W^{t-1} \setminus \{a\}) \cup \{a'\} \prec_i W^{t-1}$ by Axiom R. In either case this is not an improvement step for voter i .

Next, suppose $a \notin W^{t-1}$. We further split to subcases based on a' .

- If $a' \in W^{t-1}$ then $f(\mathbf{a}_{-i}, a') = \{a'\}$. Then $a' = x$, as otherwise $f(\mathbf{a}_{-i}, x) = \{x\} \succ_i \{a'\}$, and i is strictly better off by voting for x . This entails $W^t = \{x\}$, $R^t = W^{t-1} \setminus \{x\}$ so all invariants 1-3 hold: (1) $W^{t-1} = W^t \cup R^t$; (2) follows from (1) since $a' = x$ is the most preferred in W^{t-1} ; and (3) follows from (1) since $a = a_i^{t-1}$ is the most preferred in C' , and $a' \in C'$.
- If $a' \in R^{t-1}$ then $f(\mathbf{a}_{-i}, a') = \{a'\} \cup W^{t-1}$. Then $a' = y$, as otherwise $f(\mathbf{a}_{-i}, y) = \{y\} \cup W^{t-1} \succ_i \{a'\} \cup W^{t-1}$ by Axiom R, which means i is strictly better off by voting for y . This entails $W^t = \{y\} \cup W^{t-1}$, $R^t = R^{t-1} \setminus \{y\}$. We also get that $a' = y \succ_i x$ or else x would have been a strictly better reply. Thus all invariants 1-3 hold: (1) $W^t = W^{t-1} \cup \{y\} \subseteq W^{t-1} \cup R^{t-1}$ and $R^t = R^{t-1} \setminus \{y\}$; (2) follows from (1) since $a' = a_i^t = y$ is most preferred in R^{t-1} and strictly preferred to x ; (3) follows from (1) as in the previous case.
- If $a' \notin W^{t-1} \cup R^{t-1}$, then $W^t = f(\mathbf{a}_{-i}, a') = W^{t-1}$. The outcome does not change so this cannot be an improvement step for i .

□

Cardinal utilities A (cardinal) utility function is a mapping of candidates to real numbers $u : C \rightarrow \mathbb{R}$, where $u_i(c) \in \mathbb{R}$ is the utility of candidate c to agent i . We say that u is *consistent* with a preference relation Q_i if $u(c) > u(c') \Leftrightarrow c \succ_i c'$. The definition of cardinal utility naturally extends to multiple winners by setting $u_i(W) = \frac{1}{|W|} \sum_{c \in W} u_i(c)$ for any subset $W \subseteq C$.⁸

Lemma 15. *Consider any cardinal utility function u and the partial preference order Q it induces on subsets by random tie-breaking. Q holds Axioms K+G+R.*

The proof is rather straight-forward, and is deferred to the appendix.

Proposition 16. *f^{PR} is not FIP from the truth.*

Example 16. We use 5 candidates with initial score $\hat{s} = (1, 1, 2, 0, 0)$, and 2 agents with utilities $u_1 = (5, 3, 2, 8, 0)$ and $u_2 = (4, 2, 5, 0, 8)$. In particular, $\{b, c\} \succ_1 c$, $\{a, c\} \succ_1 \{a, b, c\}$, and $\{a, b, c\} \succ_2 \{b, c\}$, $c \succ_2 \{a, c\}$, and the following cycle occurs: $(1, 1, 2, 1, 1)\{c\} \xrightarrow{1} (1, 2, 2, 0, 1)\{b, c\} \xrightarrow{2} (2, 2, 2, 0, 0)\{a, b, c\} \xrightarrow{1} (2, 1, 2, 1, 0)\{a, c\} \xrightarrow{2} (1, 1, 2, 1, 1)\{c\}$. ◇

Finally, in contrast to the lexicographic case, convergence is no longer guaranteed if agents start from an arbitrary profile of votes, or are allowed to use direct-replies that are not best-replies. The following example shows that in the randomized tie-breaking setting even direct best reply dynamics may have cycles, albeit for specific utility scales.

⁸One interpretation is that we randomize the final winner from the set W , and hence the term randomized tie-breaking. For a thorough discussion of cardinal and ordinal utilities in normal form games, see [Borgers, 1993].

Proposition 17. f^{PR} is not restricted-FIP.

Example 17. There are 4 candidates $\{a, b, c, x\}$ and 3 agents with utilities $u_1 = (7, 3, 0, 4)$, $u_2 = (0, 7, 3, 4)$ and $u_3 = (3, 0, 7, 4)$. In particular, the following preference relations hold: $a \succ_1 \{a, b\} \succ_1 x \succ_1 \{a, c\}$; $b \succ_2 \{b, c\} \succ_2 x \succ_2 \{a, b\}$; and $c \succ_3 \{a, c\} \succ_3 x \succ_3 \{b, c\}$.

Consider the initial state $\mathbf{a}_0 = (a, b, x)$ with $s(\mathbf{a}_0) = (1, 1, 0, 1)$ and the outcome $\{a, b, x\}$. We have the following cycle where every step is the unique reply of the playing agent.

$$\begin{array}{ccccc}
 (1, 1, 0, 1)\{a, b, x\} & \xrightarrow{2} & (1, 0, 0, 2)\{x\} & \xrightarrow{3} & (1, 0, 1, 1)\{a, x, c\} \\
 \uparrow 1 & & & & \downarrow 1 \\
 (0, 1, 0, 2)\{x\} & \xleftarrow{3} & (0, 1, 1, 1)\{x, b, c\} & \xleftarrow{2} & (0, 0, 1, 2)\{x\}
 \end{array}$$

◇

Proposition 18. f^{PR} is not FDRP even from the truth.

Example 18. We take the game from Ex. 17, and add for each voter $i \in \{1, 2, 3\}$ a candidate d_i , s.t. $u_i(d_i) = 8$, $u_i(d_j) = j$ for $j \neq i$. We also add an initial score of 3 to each of the candidates $\{a, b, c, x\}$. Voter 3 moves first to $a_3^1 = x$, which is a direct reply. Then voters 1 and 2 move to their best replies a, b , respectively. Now the cycle continues as in Ex. 17. ◇

4.4 Stochastic Dominance and Local Dominance

While assigning cardinal utilities is one way to deal with ties, it is sometimes preferable not to assume a particular cardinal utility scale. Denote by $f^P(\mathbf{a}) \subseteq C$ the subset of candidates with maximal Plurality score, before any tie-breaking takes place. We can still derive a well-defined dynamics from any partial order over subsets of candidates, by assuming that a voter performs a better-response step if she strictly prefer the new outcome, and otherwise (if the new outcome is same, worse, or incomparable) she does not move.

One example of such a partial order is *stochastic dominance* (SD), which was applied to tie-breaking by [Reyhani and Wilson, 2012]. A different partial order is implied by *local dominance* (LD) which was defined for voting with uncertainty about the outcome [Conitzer *et al.*, 2011; Meir *et al.*, 2014], when uncertainty is regarding the tie breaking. We show how convergence results for LD/SD dynamics fit with other results.

Stochastic dominance Reyhani and Wilson assume that ties are broken uniformly at random, and that a voter will only perform a step that stochastically dominates the current winner(s), if such exists.

Theorem 19 (Theorem 5.7 in [Reyhani and Wilson, 2012]). *Plurality with stochastic dominance tie-breaking is FDBRP.*

We can show the following (see appendix):

Proposition 20. *A step $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is a better-response under random tie-breaking and stochastic dominance, if and only if $f^P(\mathbf{a}') \succ_i f^P(\mathbf{a})$ is entailed by Q_i , Axioms K+G+R, and transitivity.*

In other words, while Theorem 14 allowed any moves *consistent* with the axioms, SD allows only moves that *follow* from the axioms, and explicitly forbid any other step. Thus it is more restricted than expected-utility based randomized tie-breaking.

Since any SD step is also a better-reply under any cardinal utility scale, any strong or restricted convergence result for the latter applies to the former, but not vice-versa.

Local dominance Suppose that there are several candidates with maximal score. A voter may consider all of them as “perhaps winners,” without specifying how the actual winner is selected. If the voter is concerned about making a move that will leave her worse off, she will only make moves that will improve her utility with certainty, i.e. that dominates her current action (where possible worlds are all strict tie-breaking orders) [Conitzer *et al.*, 2011; Meir *et al.*, 2014; Meir, 2015].⁹

Theorem 21 (Theorem 11 in the full version of [Meir, 2015]). *Plurality with Local-Dominance tie-breaking is FDRP.*

To see how this compares with other convergence results, we need the following proposition (see appendix).

Proposition 22. *A step $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is a better-response under unknown tie-breaking and local dominance, if and only if $f^P(\mathbf{a}') \succ_i f^P(\mathbf{a})$ is entailed by Q_i , Axioms K+G, and transitivity.*

Note that since Axioms K+G+R include K+G, any LD step is also an SD step, so a restriction to LD can only eliminate cycles. Thus FBDRP follows from Theorem 19. We note that with either SD or LD tie-breaking there may be new stable states that are not Nash-equilibria. Even so, an analysis of Ex. 13 shows that all steps are entailed by Axioms K+G (and thus by Axioms K+G+R). Thus neither game form is FIP.

What if we assume that voters are even more risk-averse and only follow steps that are better-replies by Axiom K? Then it is easy to see that only moves to a more-preferred candidate can be better-replies (any move to or from a tie cannot follow from Axiom K and is thus forbidden), which means there are trivially no cycles.

5 Weak Acyclicity

Except for Plurality and Veto, convergence is not guaranteed even under restrictions on the action scheduler and the initial state. In contrast, simulations [Grandi *et al.*, 2013; Meir *et al.*, 2014; Koolyk *et al.*, 2016] show that iterative voting almost always converges even when this is not guaranteed by theory. We believe that weak acyclicity is an important part of the explanation to this gap.

5.1 Plurality with Random tie-breaking

We have seen in Section 4 that while f^{PR} is FDRP from the truthful initial state, this is no longer true from arbitrary states, and in fact f^{PR} is not restricted-FIP under any action scheduler. Our main theorem in this section shows that under a certain scheduler (of agents+actions), convergence is guaranteed from *any* state. Further, this still holds if actions are restricted to direct-replies.

Lemma 23. *Consider any game $G = \langle f_s^{PR}, \mathbf{Q} \rangle$. Consider some candidate a^* , and suppose that in \mathbf{a}^0 , there are x, y s.t. $s^0(x) \geq s^0(y) \geq s^0(a^*) + 2$. Then for any sequence of direct replies, $a^* \notin f(\mathbf{a}^t)$.*

Proof. We show that at any time $t \geq 0$ there are x^t, y^t s.t. $s^0(x), s^0(y) \geq s^0(a^*) + 2$. For $t = 0$ this holds for $x^t = x, y^t = y$. Assume by induction that the premise holds for \mathbf{a}^{t-1} . Then there are two cases:

1. $|f(\mathbf{a}^{t-1})| \geq 2$. Then since step t must be a direct reply, it must be to some candidate z with $s^{t-1}(z) \geq sw^{t-1} - 1$. Also, either x^{t-1} or y^{t-1} did not lose votes (w.l.o.g. x^{t-1}). Thus $s^t(x), s^t(z) \geq sw^{t-1} \geq s^{t-1}(a^*) + 2 \geq s^t(a^*) + 2$.
2. $|f(\mathbf{a}^{t-1})| = 1$. Then suppose $f(\mathbf{a}^{t-1}) = \{x^{t-1}\}$, and we have that $sw^{t-1} \geq s^{t-1}(a^*) + 3$. The next step is z where either $s^{t-1}(z) = sw^{t-1} - 1$ (and then we conclude as in case 1), or $s^{t-1}(z) = sw^{t-1} - 2$ and x^{t-1} loses 1 vote. In the latter case, $s^t(x^{t-1}) = s^t(z) = sw^{t-1} - 1 \geq s^{t-1}(a^*) + 2 \geq s^t(a^*) + 2$. \square

Theorem 24. f_s^{PR} is weak-FDRP.

⁹[Meir *et al.*, 2014; Meir, 2015] consider more general uncertainty over candidates' score, and [Conitzer *et al.*, 2011] considers arbitrary information sets.

Proof. Consider a game $G = \langle f_s^{PR}, \mathbf{Q} \rangle$, and an initial state \mathbf{a}^0 . For a state \mathbf{a} , denote by $B(\mathbf{a}) \subseteq A^n$ all states reachable from \mathbf{a} via paths of direct replies. Let $B = B(\mathbf{a}^0)$, and assume towards a contradiction that B does not contain a Nash equilibrium. For every $\mathbf{b} \in B$, let $C(\mathbf{b}) = \{c \in C : \exists \mathbf{a} \in B(\mathbf{b}) \wedge c \in f(\mathbf{a})\}$, i.e. all candidates that are winners in some state reachable from \mathbf{b} .

For any $\mathbf{b} \in B(\mathbf{a}^0)$, define a game $G_{\mathbf{b}}$ by taking G and eliminating all candidates *not in* $C(\mathbf{b})$. Since we only consider direct replies, for any $\mathbf{a} \in B(\mathbf{b})$, the set of outgoing edges $I(\mathbf{a})$ is the same in G and in $G_{\mathbf{b}}$ (as any direct reply must be to candidate in $C(\mathbf{b})$). Thus by our assumption, the set $B(\mathbf{b})$ in game $G_{\mathbf{b}}$ does not contain an NE.

For any $\mathbf{b} \in B(\mathbf{a}^0)$, let \mathbf{b}^* be the truthful state of game $G_{\mathbf{b}}$, and let $T(\mathbf{b}) \subseteq N$ be the set of agents who are truthful in \mathbf{b} . That is, $i \in T(\mathbf{b})$ if $b_i = b_i^*$.

Let \mathbf{b}^0 be some state $\mathbf{b} \in B(\mathbf{a}^0)$ s.t. $|T(\mathbf{b})|$ is maximal, and let $T^0 = T(\mathbf{b}^0)$. If $|T^0| = n$ then \mathbf{b}^0 is the truthful state of $G_{\mathbf{b}^0}$, and thus by Theorem 14 all best-reply paths from \mathbf{b}^0 in $G_{\mathbf{b}^0}$ lead to an NE, in contradiction to $B(\mathbf{b}^0)$ not containing any NE. Thus $T^0 < n$. We will prove that there is a path from \mathbf{b}^0 to a state \mathbf{b}' s.t. $|T(\mathbf{b}')| > |T^0|$.

Let $i \notin T(\mathbf{b}^0)$ (must exist by the previous paragraph). Consider the score of candidate b_i^* at state \mathbf{b}^0 . We divide into 5 cases. All scores specified below are in the game $G_{\mathbf{b}^0}$.

- Case 1. $|f(\mathbf{b}^0)| > 1$ and $b_i^* \in f(\mathbf{b}^0)$ (i.e. b_i^* is one of several winners). Then consider the step $\mathbf{b}^0 \xrightarrow{i} b_i^*$. This makes b_i^* the unique winner, and thus it is a direct best-reply for i . In the new state $\mathbf{b}' = (\mathbf{b}_{-i}^0, b_i^*)$ we have $T(\mathbf{b}') = T(\mathbf{b}^0) \cup \{i\}$.
- Case 2. $s^0(b_i^*) = sw^0 - 1$ (i.e., b_i^* needs one more vote to become a winner). By Axioms G+R, i prefers $f(\mathbf{b}_{-i}^0, b_i^*)$ over $f(\mathbf{b}^0)$. Then similarly to case 1, i has a direct step $\mathbf{b}^0 \xrightarrow{i} b_i^*$, which results in a “more truthful” state \mathbf{b}' .
- Case 3. $b_i^* = f(\mathbf{b}^0)$ (i.e. b_i^* is the unique winner). Then the next step $\mathbf{b}^0 \xrightarrow{j} \mathbf{b}^1$ will bring us to one of the two previous cases. Moreover, it must hold that $j \notin T(\mathbf{b}^0)$ since otherwise $b_j^0 = b_j^* = f(\mathbf{b}^0)$ which means $I_j(\mathbf{b}^0) = \emptyset$. Thus $|T(\mathbf{b}')| = |T(\mathbf{b}^1)| + 1 \geq |T(\mathbf{b}^0)| + 1$.
- Case 4. $f(\mathbf{b}^0) = x \neq b_i^*$, and $s^0(x) = s^0(b_i^*) + 2$. We further divide into:
 - Case 4.1. $s^0(b_i^*) \geq s^0(y)$ for all $y \neq x$. Then the next step by j must be from x , which brings us to one of the two first cases (as in Case 3).
 - Case 4.2. There is $y \neq x$ s.t. $s^0(x) = s^0(y) + 1 = s^0(b_i^*) + 2$. Then we continue the sequence of steps until the winner's score decreases. Since all steps that maintain sw^t select a more preferred candidate, this most occur at some time t , and $T(\mathbf{b}^0) \subseteq T(\mathbf{b}^t)$. Then at \mathbf{b}^t we are again in Case 1 or 2.
 - Case 4.3. There is $y \neq x$ s.t. $s^0(x) = s^0(y) = s^0(b_i^*) + 2$. Then by Lemma 23 b_i^* can never be selected, in contradiction to $b_i^* \in C(\mathbf{b}^0)$.
- Case 5. $f(\mathbf{b}^0) = x \neq b_i^*$, and $s^0(x) \geq s^0(b_i^*) + 3$. We further divide into:
 - Case 5.1. For all $y \neq x$, $s^0(y) \leq s^0(x) - 3$. In this case no reply is possible.
 - Case 5.2. There is some $y \neq x$ s.t. $s^0(y) \geq s^0(b_i^*) + 2$. Then by Lemma 23 b_i^* can never be selected, in contradiction to $b_i^* \in C(\mathbf{b}^0)$.
 - Case 5.3. There is some $y \neq x$ s.t. $s^0(y) \geq s^0(b_i^*) + 1$. Then the next step must be from x to such y . Which means $s^1(x) = s^1(y) = sw^0 - 1 \geq s^0(b_i^*) + 2 = s^1(b_i^*) + 2$. Thus again by Lemma 23 we reach a contradiction.

Therefore we either construct a path of direct replies to $\mathbf{b}' \in B(\mathbf{b}^0)$ with $|T(\mathbf{b}')| > |T(\mathbf{b}^0)|$ in contradiction to our maximality assumption, or we reach another contradiction. Thus $B(\mathbf{b}^0)$ must contain some NE (both in $G_{\mathbf{b}^0}$ and in G), which means by construction that G is weakly-FDRP from \mathbf{b}^0 . However since $\mathbf{b}^0 \in B(\mathbf{a}^0)$, we get that G is weakly-FDRP from \mathbf{a}^0 as well. \square

Remark 5.1. Theorem 24 and Ex. 17 provide a partial answer to an open question regarding whether there are game forms that admit weak FIP but not restricted FIP [Kukushkin, 2011]. Indeed, the game form f^{PR} for $m = 4, n = 3$ is such an example, but one that uses randomization. However if we think of f^{PR} as a deterministic game form with $2^m - 1$ possible outcomes (all nonempty subsets of candidates), where players are restricted to m actions each, then the allowed utility profiles are constrained (by Axioms G and R) and thus this result does not settle Kukushkin’s question completely.

5.2 Weighted Plurality

When voters are weighted, cycles of direct responses can emerge [Meir *et al.*, 2010; Meir, 2016]. We conjecture that such cycles must depend on the order of agents, and that certain orders will break such cycles and reach an equilibrium, at least from the truthful state.

Conjecture 25. $f_{\hat{s}, \mathbf{w}}^{PL}$ is weak-FDRP (in particular weak-FIP).

Similar techniques to those used so far appear to be insufficient to prove the conjecture. For example, in contrast to the unweighted case, a voter might return to a candidate she deserted in *any scheduler*, even if only two weight levels are present. We thus leave the proof of the general conjecture for future work.

Yet, we want to demonstrate the power of weak acyclicity over restricted acyclicity, even when there are no randomness or restrictions on the utility space. That is, to provide a definite (negative) answer to Kukushkin’s question of whether weak acyclicity entails restricted acyclicity. To do so, we will use a slight variation of Plurality with weighted voters and lexicographic tie-breaking.

Theorem 26. There exist a game form f^* s.t. f^* is weak-FIP but not restricted-FIP.

Proof. Consider the following game G : The initial fixed score of candidates $\{a, b, c, d\}$ is $\hat{s} = (0, 1, 2, 3)$. The weight of each voter $i \in \{1, 2, 3\}$ is i . The preference profile is as follows: $c \succ_1 d \succ_1 b \succ_1 a$, $b \succ_2 c \succ_2 a \succ_2 d$, and $a \succ_3 b \succ_3 c \succ_3 d$. This game was used in [Meir *et al.*, 2010] to demonstrate that Plurality with weighted voters is not FDRP, however it can be verified that G is restricted-FIP so it is not good enough for our use.

If we ignore agents’ preferences, we get a particular game form $f_{\hat{s}, \mathbf{w}}^{PL}$ where $N = \{1, 2, 3\}$, $M = \{a, b, c, d\}$, $\hat{s} = (0, 1, 2, 3)$ and $\mathbf{w} = (1, 2, 3)$.

We define f^* by modifying $f_{\hat{s}, \mathbf{w}}^{PL}$ with the following restrictions on agents’ actions: $A_1 = \{c, d\}$, $A_2 = \{b, c\}$, $A_3 = \{a, b, d\}$. Thus f^* is a $2 \times 2 \times 3$ game form, presented in Figure 6(a).

We first show that f^* is not restricted-FIP. Indeed, consider the game G^* accepted from f^* with the same preferences from game G (Figure 6(b)). We can see that there is a cycle of length 6 (in bold). An agent scheduler that always selects the agent with the bold reply guarantees that convergence does not occur, since in all 6 relevant states the selected agent has no alternative replies.

Next, we show that f^* is weak-FIP. That is, for any preference profile there is some scheduler that guarantees convergence. We thus divide into cases according to the preferences of agent 3. In each case, we specify a state where the scheduler selects agent 3, the action of the agent, and the new state.

We note that since all thick edges must be oriented in the same direction, $a \succ_3 b$ if and only if $b \succ_3 c$. Thus the following three cases are exhaustive.

	Q_3	state	action	new state
1	$b \succ d$	(d, b, a)	b	(d, b, b)
2	$d \succ b \ \& \ d \succ a$	(c, b, b)	d	(c, b, d)
3	$a \succ d \succ b \succ c$	(d, c, b)	d	(d, c, d)

In either case, agent 3 moves from a state on the cycle to a Nash equilibrium. □

6 Conclusions and Future Work

The main conceptual contribution of this work was to provide a joint rigorous framework for the study of iterative voting, as part of the broader literature on acyclicity of games and game forms.

Voting rule	FIP	FBRP	FDBRP	restricted-FIP	Weak-FIP
Dictator	V	V	-	V	V
Plurality (lex.)	X	X (Ex. 6)	V (Thm. 5)	V	V
Plurality (LD)	X (Ex. 13)	?	V [M15]	V	V
Plurality (SD)	X (Ex. 13)	?	V [RW12]	V	V
Plurality (rand.)	X (Ex. 13)	X	X	X (Ex. 17)	V (Thm. 24)
Weighted Plurality (lex.)	X	X	X (Ex. 7)	?	?
Veto	X	X [M16]	V [RW12,LR12]	V	V
k -approval ($k \geq 2$)	X	X [LR12,L15]	-	X	X [M16]
Borda	X	X [RW12,LR12]	-	X	X [RW12]
PSRs (except k -approval)	X	X [LR12,L15]	-	?	?
Approval	X	X [M16]	-	V [M16]	V
Other common rules	X	X [KLR16]	-	?	?

Table 1: Positive results carry to the right side, negative to the left side. We assume lexicographic tie breaking in all rules except Plurality. FDBRP is only well-defined for Plurality and Veto. Reference codes: RW12 [Reyhani and Wilson, 2012], LR12 [Lev and Rosenschein, 2012], M15 [Meir, 2015], L15 [Lev, 2015], M16 [Meir, 2016], KLR16 [Koolyk *et al.*, 2016].

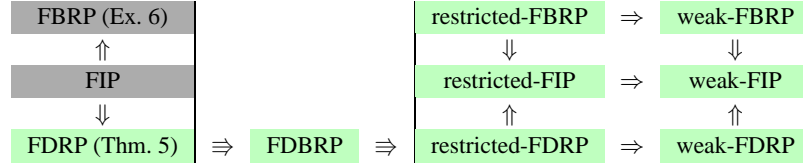


Figure 7: Convergence results for Plurality under lexicographic tie-breaking. Positive results (in light green) carry with the direction of the arrows, whereas negative results (dark gray) carry in the opposite direction.

(at least) all unanimous votes form equilibria. Another sufficient condition due to Apt and Simon [2012] is by eliminating never-best-reply strategies, and the prospects of applying it to common voting rules is not yet clear.

We can see that in the “standard” lexicographic domain, convergence is guaranteed from any initial state provided that voters restrict themselves to direct replies. With randomized tie-breaking, we must also require a truthful initial vote. On the other hand, we can also allow indirect best-replies, so the results are essentially incomparable. However, we see the result in the lexicographic case as stronger, since it only requires a very mild and natural behavioral restriction in the context of Plurality voting, whereas it is harder to justify assumptions on the initial state.

Implications on social choice Importantly, best-reply dynamics is a natural and straightforward process, and requires little information. As such, and due to the convergence properties demonstrated in this work, it is an attractive “baseline” candidate for predicting human voter behavior in elections and designing artificial agents with strategic voting capabilities—two of the most important, and also the hardest, goals of social choice research. However, the clear disadvantage of this approach is that in the vast majority of cases (especially when there are more than a handful of voters), almost every voting profile (including the truthful one) is already a Nash equilibrium. Given this, our analysis is particularly suitable when the number of voters is small, for two main reasons. First, it is more practical to perform an iterative voting procedure with few participants. Second, the question of convergence is only relevant when cases of tie or near-tie are common. In more complex situations with many active voters who may change their vote, it is likely that a more elaborate game-theoretic model is required, which takes into account voters’ uncertainty and heuristic behavior (see Section 1.1).




Figure 8: Convergence results for Plurality under random tie-breaking.

Promising future directions Based on the progress made in this paper and the other results published since the introduction of iterative voting in [Meir *et al.*, 2010], we believe that research in this area should focus on three primary directions:

1. Weak-acyclicity seems more indicative than order-free acyclicity to determine convergence in practice. Thus theorists should study which voting rules are weak-FIP, perhaps under reasonable restrictions (as we demonstrated, this property is distinct from restricted-FIP). We highlight that even in rules where there are counter examples for weak acyclicity (k-approval, Borda), these examples use two voters and games with more voters may well be weakly acyclic.
2. It is important to experimentally study how people really vote in iterative settings (both in and out of the lab), so that this behavior can be formalized and behavioral models can be improved. The work of [Tal *et al.*, 2015] is a preliminary step in this direction, but there is much more to learn. Ideally, we would like to identify a few types of voters, such that for each type we can relatively accurately predict the next action in a particular state. It would be even better if these types are not specific to a particular voting rule or contextual details.
3. We would like to know not only if a voting rule converges under a particular dynamics (always or often), but also what are the properties of the attained outcome—in particular, whether the iterative process improves welfare or fairness, avoids “voting paradoxes” [Xia *et al.*, 2007] and so on. Towards this end, several researchers (e.g., [Reijngoud and Endriss, 2012; Brânzei *et al.*, 2013; Meir *et al.*, 2014; Bowman *et al.*, 2014; Koolyk *et al.*, 2016]) have started to explore these questions via theory and simulations. However, a good understanding of how iterative voting shapes the outcome, whether the population of voters consists of humans or artificial agents, is still under way.

References

- Stéphane Airiau and Ulle Endriss. Iterated majority voting. In *Proceedings of the 1st International Conference on Algorithmic Decision Theory (ADT-2009)*, pages 38–49. Springer Verlag, 2009.
- Daniel Andersson, Vladimir Gurvich, and Thomas Dueholm Hansen. On acyclicity of games with cycles. *Discrete Applied Mathematics*, 158(10):1049–1063, 2010.
- Krzysztof R. Apt and Sunil Simon. A classification of weakly acyclic games. In *SAGT’12*, pages 1–12. 2012.
- Tilman Borgers. Pure strategy dominance. *Econometrica*, 61(2):423–430, 1993.
- Endre Boros, Vladimir Gurvich, Kazuhisa Makino, and Wei Shao. Nash-solvable bidirected cyclic two-person game forms. Technical report, Rutcor Research Report 26-2007 and DIMACS Technical Report 2008-13, Rutgers University, 2008.
- Endre Boros, Vladimir Gurvich, Kazuhisa Makino, and Dávid Papp. Acyclic, or totally tight, two-person game forms: Characterization and main properties. *Discrete Mathematics*, 310(6):1135–1151, 2010.

- Michael Bowling. Convergence and no-regret in multiagent learning. *Advances in neural information processing systems*, 17:209–216, 2005.
- Clark Bowman, Jonathan K. Hodge, and Ada Yu. The potential of iterative voting to solve the separability problem in referendum elections. *Theory and decision*, 77(1):111–124, 2014.
- Simina Brânzei, Ioannis Caragiannis, Jamie Morgenstern, and Ariel D. Procaccia. How bad is selfish voting? In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- Samir Chopra, Eric Pacuit, and Rohit Parikh. Knowledge-theoretic properties of strategic voting. In *Logics in Artificial Intelligence*, pages 18–30. Springer, 2004.
- Vincent Conitzer, Toby Walsh, and Lirong Xia. Dominating manipulations in voting with partial information. In *AAAI*, volume 11, pages 638–643, 2011.
- Antoine-Augustin Cournot. *Recherches sur les principes mathématiques de la théorie des richesses par Augustin Cournot*. chez L. Hachette, 1838.
- Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proc. of 11th ACM-EC*, pages 347–356, 2010.
- Amrita Dhillon and Ben Lockwood. When are plurality rule voting games dominance-solvable? *Games and Economic Behavior*, 46:55–75, 2004.
- Edith Elkind, Umberto Grandi, Francesca Rossi, and Arkadii Slinko. Gibbard–satterthwaite games. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-15)*, 2015.
- Ulle Endriss. Sincerity and manipulation under approval voting. *Theory and Decision*, 74(3):335–355, 2013.
- Alex Fabrikant, Aaron D Jagard, and Michael Schapira. On the structure of weakly acyclic games. In *Algorithmic Game Theory*, pages 126–137. Springer, 2010.
- Dvir Falik, Reshef Meir, and Moshe Tennenholtz. On coalitions and stable winners in plurality. In *WINE’12*, pages 256–269. 2012.
- Robert Forsythe, Thomas Rietz, Roger Myerson, and Robert Weber. An experimental study of voting rules and polls in three-candidate elections. *International Journal of Game Theory*, 25(3):355–83, 1996.
- Ehud Friedgut, Gil Kalai, Nathan Keller, and Noam Nisan. A quantitative version of the gibbard-satterthwaite theorem for three alternatives. *SIAM Journal on Computing*, 40(3):934–952, 2011.
- Peter Gärdenfors. Manipulation of social choice functions. *Journal of Economic Theory*, 13(2):217–228, 1976.
- Christian Geist and Ulle Endriss. Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial Intelligence Research*, 40(1):143–174, 2011.
- A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
- Neelam Gohar. *Manipulative Voting Dynamics*. PhD thesis, University of Liverpool, 2012.
- Umberto Grandi, Andrea Loreggia, Francesca Rossi, Kristen Brent Venable, and Toby Walsh. Restricted manipulation in iterative voting: Condorcet efficiency and Borda score. In *Algorithmic Decision Theory*, pages 181–192. Springer, 2013.
- Richard W Hamming. Error detecting and error correcting codes. *Bell System technical journal*, 29(2):147–160, 1950.

- Jerry S Kelly. Strategy-proofness and social choice functions without singlevaluedness. *Econometrica: Journal of the Econometric Society*, pages 439–446, 1977.
- Aaron Koolyk, Omer Lev, and Jeffrey S Rosenschein. Convergence and quality of iterative voting under non-scoring rules. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 1329–1330. International Foundation for Autonomous Agents and Multiagent Systems, 2016.
- Nikolai S. Kukushkin. Congestion games: a purely ordinal approach. *Economics Letters*, 64:279–283, 1999.
- Nikolai S. Kukushkin. Perfect information and congestion games. *Games and Economic Behavior*, 38:306–317, 2002.
- Nikolai S. Kukushkin. Acyclicity of improvements in finite game forms. *International Journal of Game Theory*, 40(1):147–177, 2011.
- Omer Lev and Jeffrey S. Rosenschein. Convergence of iterative voting. In *Proc. of 11th AAMAS*, pages 611–618, 2012.
- Omer Lev. *Agent Modeling of Human Interaction: Stability, Dynamics and Cooperation*. PhD thesis, The Hebrew University of Jerusalem, 2015.
- Jason R Marden, Gürdal Arslan, and Jeff S Shamma. Regret based dynamics: convergence in weakly acyclic games. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, page 42. ACM, 2007.
- Reshef Meir, Maria Polukarov, Jeffrey S. Rosenschein, and Nick Jennings. Convergence to equilibria of plurality voting. In *Proc. of 24th AAAI*, pages 823–828, 2010.
- Reshef Meir, Omer Lev, and Jeffrey S. Rosenschein. A local-dominance theory of voting equilibria. In *EC’14*, 2014.
- Reshef Meir. Plurality voting under uncertainty. In *AAAI’15*, pages 2103–2109, 2015.
- Reshef Meir. Strong and weak acyclicity in iterative voting. In *COMSOC’16*. 2016.
- Matthias Messner and Mattias K. Polborn. Robust political equilibria under plurality and runoff rule, 2002. Mimeo, Bocconi University.
- Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13(1):111–124, 1996.
- Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- Svetlana Obraztsova, Evangelos Markakis, and David R. M. Thompson. Plurality voting with truth-biased agents. In *Algorithmic Game Theory*, pages 26–37. Springer, 2013.
- Svetlana Obraztsova, Evangelos Markakis, Maria Polukarov, Zinovi Rabinovich, and Nicholas R Jennings. On the convergence of iterative voting: How restrictive should restricted dynamics be? In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Zinovi Rabinovich, Svetlana Obraztsova, Omer Lev, Evangelos Markakis, and Jeffrey S. Rosenschein. Analysis of equilibria in iterative voting schemes. In *29th National Conference on AI (AAAI)*, 2015.
- Annemiek Reijngoud and Ulle Endriss. Voter response to iterated poll information. In *Proc. of 11th AAMAS*, pages 635–644, 2012.
- Reyhaneh Reyhani and Mark C Wilson. Best-reply dynamics for scoring rules. In *20th European Conference on Artificial Intelligence*. IOS Press, 2012.

- Alvin E Roth. The college admissions problem is not equivalent to the marriage problem. *Journal of economic Theory*, 36(2):277–288, 1985.
- Donald G. Saari. Susceptibility to manipulation. *Public Choice*, 64:21–41, 1990.
- M. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- Murat R. Sertel and M. Remzi Sanver. Strong equilibrium outcomes of voting games are the generalized condorcet winners. *Social Choice and Welfare*, 22:331–347, 2004.
- Maor Tal, Reshef Meir, and Ya’acov Gal. A study of human behavior in voting systems. In *AAMAS 15*, pages 665–673, 2015.
- Lirong Xia, Jérôme Lang, and Mingsheng Ying. Sequential voting rules and multiple elections paradoxes. In *TARK’07*, pages 279–288, 2007.
- H Peyton Young. The evolution of conventions. *Econometrica: Journal of the Econometric Society*, pages 57–84, 1993.

A Proofs

Lemma 15. Consider any cardinal utility function u and the partial preference order Q it induces on subsets by random tie-breaking. Q holds Axioms $K+G+R$.

Proof. Let u be any utility scale, we will show that all axioms hold. Let $a, b \in C$ and $W \subseteq C \setminus \{a, b\}$.

$$u(\{a\} \cup W) = \frac{1}{|W|+1} \left(u(a) + \sum_{c \in W} u(c) \right), u(\{b\} \cup W) = \frac{1}{|W|+1} \left(u(b) + \sum_{c \in W} u(c) \right) = u(\{b\} \cup W),$$

thus $\{a\} \cup W \succ_Q \{b\} \cup W$, and Axiom R holds.

Let $a \in C, W \subseteq C$ s.t. $\forall b \in W, a \succ_b$. Then

$$\begin{aligned} u(a) &= \frac{1}{|W|+1} \left(u(a) + \sum_{b \in W} u(a) \right) > \frac{1}{|W|+1} \left(u(a) + \sum_{b \in W} u(b) \right) = u(\{a\} \cup W) \\ &> \frac{1}{|W|+1} \left(u(W) + \sum_{b \in W} u(a) \right) = \frac{1}{|W|+1} u(W) + \frac{|W|}{|W|+1} u(W) = u(W), \end{aligned}$$

thus $a \succ_Q \{a\} \cup W \succ_Q W$ and Axiom G holds.

Axiom K1 follows immediately from G. K2 also follows if preferences are strict. Even if there are ties, and $a \succeq w$ for all $a \in A, w \in W$ then:

$$u(A) \geq \min_{a \in A} u(a) \geq \max_{w \in W} u(w) \geq u(W),$$

i.e., $A \succeq_Q W$. □

Definition A.1. Suppose that $X, Y \subseteq C, k = |X| \leq |Y| = K$. Sort X, Y in increasing order by Q . Let $r_j = \lceil \frac{j}{k} K \rceil$. Partition Y into sets Y_1, \dots, Y_k s.t. for $j < K, Y_j = \{y_{r_{j-1}+1}, \dots, y_{r_j}\}$ (e.g., if $k = 3, K = 7$, then Y is partitioned into $Y_1 = \{y_1, y_2, y_3\}, Y_2 = \{y_4, y_5\}, Y_3 = \{y_6, y_7\}$).

X match-dominates Y according to Q if:

- (I) $\forall j \leq k \forall y \in Y_j, x_j \succeq y$; and
- either (IIa) at least one relation is strict, or (IIb) $K \bmod k \neq 0$.

If $|X| > |Y|$, then X match-dominates Y if Y match-dominates X according to the reverse of Q .

Intuitively, match-domination means that for any $q \in [0, 1]$, there is a fraction q of the set X that dominates a fraction of $1 - q$ from the set Y : at least one $x \in X$ dominates all of Y , at least 20% of X dominate at least 80% of Y , and so on.

Lemma 27. Let \mathbf{a}, \mathbf{a}' be two profiles that differ by a single vote, and define $X = f(\mathbf{a}), Y = f(\mathbf{a}')$.¹⁰

The following conditions are equivalent for any strict order Q over C :

1. X stochastically dominates Y under preferences Q and uniform lottery.
2. The relation $X \succ Y$ is entailed by Q and the Axioms $K+G+R$ and transitivity.
3. $u(X) > u(Y)$ for every u that is consistent with Q .
4. X match-dominates Y according to Q .

¹⁰Without some restriction on X, Y , the lemma is incorrect. E.g. if $x_1 \succ y_1 \succ y_2 \succ x_2 \succ y_3 \succ y_4$, then X stochastically dominates Y but there is no way to derive $X \succ Y$ from the axioms $K+G+R$.

Proof. The equivalence of (1) and (3) is immediate, and used e.g. in [Reyhani and Wilson, 2012].

(2) \Rightarrow (3). If $X \succ Y$ follows from the axioms, then there is a sequence of sets $X = X_0 \succ X_1 \succ \dots \succ X_k = Y$ such that each $X_j \succ X_{j+1}$ follows from a single axiom K, G, or R. Thus it is sufficient to show for $X \succ Y$ that follows from a single axiom.

If $X \succ Y$ follows from Axiom R, then $X = \{a\} \cup W, Y = \{b\} \cup W$ for some $W \subseteq C \setminus \{a, b\}$ and $a \succ b$. Thus

$$u(X) = u(\{a\} \cup W) = \frac{1}{|W| + 1} \left(u(a) + \sum_{c \in W} u(c) \right) > \frac{1}{|W| + 1} \left(u(b) + \sum_{c \in W} u(c) \right) = u(\{b\} \cup W) = u(Y).$$

If $X \succ Y$ follows from Axiom G, then either $X = Y \cup \{a\}$ and $a \succ b$ for all $b \in Y$, or $X = \{x\}$ and $Y = \{x\} \cup W$ where $x \succ w$ for all $w \in W$. For the first case

$$\begin{aligned} u(X) &= \frac{1}{|Y| + 1} u(a) + \frac{1}{|Y| + 1} \sum_{y \in Y} u(y) = \frac{1}{|Y| + 1} \frac{1}{|Y|} \sum_{y \in Y} u(a) + \frac{1}{|Y| + 1} \sum_{y \in Y} u(y) \\ &> \frac{1}{|Y| + 1} \frac{1}{|Y|} \sum_{y \in Y} u(y) + \frac{1}{|Y| + 1} \sum_{y \in Y} u(y) \\ &= \left(1 + \frac{1}{|Y|} \right) \frac{1}{|Y| + 1} \sum_{y \in Y} u(y) = \frac{1}{|Y|} \sum_{y \in Y} u(y) = u(Y). \end{aligned}$$

For the second case,

$$u(X) = u(x) = \frac{1}{|Y|} \sum_{y \in Y} u(x) = \frac{1}{|Y|} \left(u(x) + \sum_{w \in W} u(x) \right) > \frac{1}{|Y|} \left(u(x) + \sum_{w \in W} u(w) \right) = u(Y).$$

If $X \succ Y$ follows from Axiom K, then $u(x) > u(y)$ for any $x \in X, y \in Y$ which is a trivial case.

(3) \Rightarrow (4). Suppose that $u(X) > u(Y)$ for all u . Suppose first $|X| \leq |Y|$. If $|X|$ does not match-dominate Y then either (I) there is an element $x_{j'}$ that is less preferred than some element $y' \in Y_{j'}$; or (II) for all j and all $y \in Y_j, x_j =_Q y$ and $|Y_j| = \frac{K}{k} = q$ for all j . We will derive a contradiction to (3) in either case. In the latter case, we have $u(x_j) = u(Y_j)$ for all j and thus

$$u(Y) = \frac{1}{K} \left(\sum_{j \leq k} |Y_j| u(Y_j) \right) = \frac{\sum_{j \leq k} q u(x_j)}{K} = \frac{\sum_{j \leq k} q u(x_j)}{kq} = u(X),$$

In contradiction to (3).

Thus we are left with case (I). That is, there are $j' \leq k$ and $y' \in Y_{j'}$ s.t. $x_{j'} \prec y'$. We define the (possibly empty) set $X' \subseteq X$ as all elements $\{x : x \succ x_{j'}\}$. We define $Y' \subseteq Y$ as $\{y : y \succeq y'\}$. By construction, for any $j > j', Y_j \subseteq Y'$. Thus

$$|Y'| \geq 1 + \sum_{j=j'+1}^k |Y_j| = 1 + \sum_{j=j'+1}^k (r_j - r_{j-1}) = (K - r_{j'}) + 1 = (K - \left\lceil \frac{j'}{k} K \right\rceil) + 1 > K - \frac{j'}{k} K = K(1 - \frac{j'}{k}),$$

whereas $|X'| \leq k - j'$. We define u as follows: $u(x) = 1, u(y) = 1$ for all $x \in X', y \in Y$, and $u(z) = 0$ for all other elements. Note that X', Y' contain the top elements of X, Y , respectively. In addition, y' is the minimal element in Y' and by transitivity $y' \succ x$ for all $x \in X \setminus X'$. Thus u is consistent with Q .

We argue that $u(Y) > u(X)$ in contradiction to (3). Indeed, $u(X) = \frac{|X'|}{|X|} \leq \frac{k-j'}{k} = 1 - \frac{j'}{k}$.

$$u(Y) = \frac{|Y'|}{|Y|} > \frac{(1 - \frac{j'}{k})K}{K} = 1 - \frac{j'}{k} = \frac{k-j'}{k} \geq \frac{|X'|}{|X|} = u(X),$$

so we get a contradiction to (3) again. Thus X matching-dominate Y .

(4) \Rightarrow (2). This is the only part of the proof where we use the profiles from which X, Y are obtained. When a single voter moves, either the winner set changes by a single candidate (added, removed, or swapped), or X is a single candidate, or Y is a single candidate. We prove case by case.

- The case where $|X| = |Y| = 1$ is immediate.
- Suppose $|X| = 1$ (i.e. $X = \{x\}$) and $|Y| = K > 1$. Then X match-dominates Y means that $x \succsim y$ for all $y \in Y$, with at least one relation being strict, w.l.o.g. y_K (least preferred in Y). Then $X \succsim \{y_1, \dots, y_{K-1}\} \succ Y$, where the first transition is by Axiom K2 and the second is by Axiom G.
- The case of $|Y| = 1$ is symmetric.
- Suppose $|X| = |Y| = k$. Then X match-dominates Y means that $x_i \succeq y_i$ for all i . For all $t \in \{0, 1, \dots, k\}$, let $X^t = \{x_1, \dots, x_t, y_{t+1}, \dots, y_k\}$. Then $X^{t-1} = X^t$ if $x_t = y_t$, and $X^{t-1} \succ X^t$ otherwise from Axiom R. In addition, $X = X^0, Y = X^k$ thus $X \succ Y$ from transitivity.
- Suppose $|X| = k, |Y| = k + 1$. Then X match-dominates Y means that $|Y_1| = \lceil \frac{k+1}{k} \rceil = 2$, and all other sets Y_j are singletons $Y_j = y_j$. Consider the set Y' that includes the top k elements of Y . Since x_1 is (weakly) preferred to both candidates in Y_1 , Y' is match-dominated by X . By the previous bullet $X \succeq Y'$ follows from Axiom R and transitivity. Finally, $Y' \succ Y = Y' \cup \{\min Y\}$ by Axiom G.

□

The following is an immediate corollary:

Proposition 20. A step $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is a better-response under random tie-breaking and stochastic dominance, if and only if $f(\mathbf{a}') \succ_i f(\mathbf{a})$ is entailed by Q_i , the Axioms K+G+R, and transitivity.

Proposition 22. A step $\mathbf{a} \xrightarrow{i} \mathbf{a}'$ is a better-response under unknown tie-breaking and local dominance, if and only if $f(\mathbf{a}') \succ_i f(\mathbf{a})$ is entailed by Q_i , Axioms K+G, and transitivity.

Proof. Suppose that $X = f(\mathbf{a}')$ locally-dominates $Y = f(\mathbf{a})$. Let $Z = X \cap Y$, and $X' = X \setminus Z, Y' = Y \setminus Z$. We must have $x \succ_i y$ for any $x \in X, y \in Y'$, otherwise, a tie-breaking order that selects y first and x second would make i strictly lose when moving from Y to X . Similarly, $x \succ_i y$ for any $x \in X', y \in Y$. If $Z = \emptyset$ then $X = X' \succ_i Y' = Y$ follows from Axiom K. Otherwise, by repeatedly applying Axiom G we get $X \succeq_i Z \succeq_i Y$ with at least one relation being strict.

In the other direction, since Axiom G can only be used to add elements lower (or higher) than all existing elements, it may only induce relations of the form $Z \succ Z \cup Y'$ where $z \succ y$ for all $z \in Z, y \in Y'$; or relations of the form $Z \cup X' \succ Z$ where $x \succ z$ for all $z \in Z, x \in X'$. Thus if $X \succ Y$ follows from Axiom G, they must have the form $X = Z \cup X', Y = Z \cup Y'$ where $x \succ z \succ y$ for all $x \in X', z \in Z, y \in Y'$. To see that this entails local dominance, let $x_L = L(X)$ be the first element in X according to order $L \in \pi(C)$, and likewise for Y . For any L , $x_L \succeq y_L$ (with equality iff $L(X) = L(Y) \in Z$). Further, either X' or Y' are non-empty (w.l.o.g. X'). Consider an order L' such that $L'(X) \in X'$, then $x_{L'} \succ y$ for all $y \in Y$ and in particular $x_{L'} \succ y_{L'}$. □